Network Working Group
Request for Comments: 4915
Category: Standards Track
P. Psenak
Cisco Systems
S. Mirtorabi
Force10 Networks
A. Roy
L. Nguyen
P. Pillay-Esnault
Cisco Systems
June 2007

Multi-Topology (MT) Routing in OSPF

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Abstract

This document describes an extension to Open Shortest Path First (OSPF) in order to define independent IP topologies called Multi-Topologies (MTs). The Multi-Topologies extension can be used for computing different paths for unicast traffic, multicast traffic, different classes of service based on flexible criteria, or an in-band network management topology.

An optional extension to exclude selected links from the default topology is also described.

Table of Contents

1. Introduction

OSPF uses a fixed packet format, therefore it is not easy to introduce any backward-compatible extensions. However, the OSPF specification [OSPF] introduced Type of Service (TOS) metric in an earlier specification [TOS-OSPF] in order to announce a different link cost based on TOS. TOS-based routing as described in [TOS-OSPF] was never deployed and was subsequently deprecated. [M-ISIS] describes a similar mechanism for ISIS.

We propose to reuse the TOS-based metric fields. They have been redefined and are used to advertise different topologies by advertising separate metrics for each of them.

1.1. Differences between Multi-Topology and TOS-Based Routing

Multi-Topology routing differs from [TOS-OSPF] TOS-based routing in the following ways:

1. With TOS routing [TOS-OSPF], the TOS or Diffserv Code Point (DSCP) in the IP header is mapped directly to the corresponding OSPF SPF calculation and routing table. This limits the number and definition of the topologies to the 16 TOS values specified in Section 12.3 of [TOS-OSPF]. With Multi-Topology routing, the classification of what type of traffic maps to which topology is not within the scope of this document.

2. With TOS routing [TOS-OSPF], traffic that is unreachable in the routing table associated with the corresponding TOS will revert to the TOS 0 routing table. With Multi-Topology routing, this is optional.

3. With TOS routing [TOS-OSPF], individual links or prefixes could not be excluded from a topology. If the Link State Advertisement (LSA) options T-bit was set, all links or prefixes were either advertised explicitly or defaulted to the TOS 0 metric. With Multi-Topology routing, links or prefixes that are not advertised for a specific topology do not exist in that topology.

2. Terminology

2.1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC-KEYWORDS].

2.2. Terms

We use the following terminology in this document:

Non-MT router
: Routers that do not have the MT capability.

MT router
: Routers that have MT capability as described in this document.

MT-ID
: Renamed TOS field in LSAs to represent Multi-Topology ID.

Default topology
: Topology that is built using the TOS 0 metric (default metric).

MT topology
: Topology that is built using the corresponding MT-ID metric.

MT
: Shorthand notation for MT topology.

MT#0 topology
: Representation of TOS 0 metric in MT-ID format.

Non-MT-Area
: An area that contains only non-MT routers.

MT-Area
: An area that contains both non-MT routers and MT routers, or only MT routers.

3. Base MT Functional Specifications

3.1. MT Area Boundary

Each OSPF interface belongs to a single area, and all MTs sharing that link need to belong to the same area. Therefore, the area boundaries for all MTs are the same, but each MT’s attachment to the area is independent.

3.2. Adjacency for MTs

Each interface can be configured to belong to a set of topologies. A single adjacency is formed with neighbors on the interface even if the interface is configured to participate in multiple topologies. Furthermore, adjacency formation is independent of the topologies configured on the local interface and the neighboring router.

3.3. Sending OSPF Control Packets

Sending OSPF control packets is unchanged from [OSPF]. For OSPF control packets sent to the remote end of a virtual link, the transit area path MUST be composed of links participating in the default topology and the OSPF control packets MUST be forwarded using the default topology.

3.4. Advertising MT Adjacencies and the Corresponding IP Prefixes

The TOS metric field is reused to advertise topology specific metric for links and prefixes belonging to that topology. The TOS field is redefined as MT-ID in the payload of Router, Summary, and Type-5 and Type-7 AS-external-LSAs (see Appendix B).

MT-ID metrics in LSAs SHOULD be in ascending order of MT-ID. If an MT-ID exists in an LSA or router link multiple times, the metric in the first MT-ID instance MUST be used.

When a router establishes a FULL adjacency over a link that belongs to a set of MTs, it advertises the corresponding cost for each MT-ID.

By default, all links are included in the default topology and all advertised prefixes belonging to the default topology will use the TOS 0 metric as in [OSPF].

Each MT has its own MT-ID metric field. When a link is not part of a given MT, the corresponding MT-ID metric is excluded from the LSA.

The Network-LSA does not contain any MT information since the Designated Router (DR) is shared by all MTs. Hence, there is no change to the Network-LSA.

3.4.1. Inter-Area and External Routing

In Summary-LSAs and Type-5 and Type-7 AS-external-LSAs, the TOS metric fields are redefined as MT-ID metric fields and are used to advertise prefix and router reachability in the corresponding topology.

When a router originates a Summary-LSA, or Type-5 or Type-7 AS-external-LSA that belongs to a set of MTs, it includes the corresponding cost for each MT-ID. By default, the prefix participates in the default topology and uses the TOS 0 metric for the default topology, similar to standard OSPF [OSPF].

Setting the P-bit in Type-7 AS-external-LSA is topology independent and pertains to all MT-ID advertised in the body of the LSA.

3.5. Flushing MT Information

When a certain link or prefix that existed or was reachable in a certain topology is no longer part of that topology or is unreachable in that topology, a new version of the LSA MUST be originated excluding metric information representing the link or prefix in that topology.

The MT metric in the Router-LSA can also be set to the maximum possible metric to enable the router to become a stub in a certain topology [STUB].

3.6. MT SPF Computation

By considering MT-ID metrics in the LSAs, OSPF computes multiple topologies and finds paths to IP prefixes for each MT independently. A separate SPF will be computed for each MT-ID to find independent paths to IP prefixes.

Network-LSAs are used by all topologies during the SPF computation. During the SPF for a given MT-ID, only the links and metrics for that MT-ID are considered. Entries in the Router Routing table are also MT-ID specific.

3.7. MT-ID Values

Since AS-External-LSAs use the high-order bit in the MT-ID field (E-bit) for the external metric-type, only MT-IDs in the 0 to 127 range are valid. The following MT-ID values are reserved:

0	- Reserved for advertising the metric associated with the default topology (see Section 4.2)
1	- Reserved for advertising the metric associated with the default multicast topology
2	- Reserved for IPv4 in-band management purposes
3-31	- Reserved for assignments by IANA
32-127	- Reserved for development, experimental and proprietary features [RFC3692]
128-255	- Invalid and SHOULD be ignored

3.8. Forwarding in MT

It is outside of the scope of this document to specify how the information in various topology specific forwarding structures are used during packet forwarding or how incoming packets are associated with the corresponding topology. For correct operation, both forwarding behavior and methods of associating incoming packets to a corresponding topology must be consistently applied in the network.

4. Default Topology Link Exclusion Functional Specifications

The Multi-Topologies imply that all the routers participate in the default topology. However, it can be useful to exclude some links from the default topology and reserve them for some specific classes of traffic.

The Multi-Topologies extension for the default topology link or prefix exclusion is described in the following subsections.

4.1. Exclusion of Links in the Default Topology

OSPF does not have the notion of an unreachable link. All links can have a maximum metric of 0xFFFF advertised in the Router-LSA. The link exclusion capability requires routers to ignore TOS 0 metrics in Router-LSAs in the default topology and to alternately use the MT-ID#0 metric to advertise the metric associated with the default topology. Hence, all routers within an area MUST agree on how the metric for the default topology will be advertised.

The unused T-bit is defined as the MT-bit in the option field in order to ensure that a Multi-Topology link-excluding capable router will only form an adjacency with another similarly configured router.

```
+---+---+---+---+---+---+---+---+
|DN |O  |DC |EA |NP |MC |E  |MT |
+---+---+---+---+---+---+---+---+
```

Figure 1: OSPF Option Bits

MT-bit: If DefaultExclusionCapability is enabled, the bit MUST be set in Hello packets and SHOULD be set in Database Description packet (see Section 4.2).

4.2. New Area Data Structure Parameter

We define a new parameter in the Area Data Structure:

DefaultExclusionCapability
This configurable parameter ensures that all routers in an area have this capability enabled before the default topology can be disabled on a router link in the area without causing backward-compatibility problems.

When an area data structure is created, the DefaultExclusionCapability is disabled by default.

If DefaultExclusionCapability is disabled:

- o The MT-bit MUST be cleared in Hello packets and SHOULD be cleared in Database Description packets.

- o If a link participates in a non-default topology, it is automatically included in the default topology to support backward compatibility between MT and non-MT routers. This is accomplished using the TOS 0 metric field as in [OSPF].

If DefaultExclusionCapability is enabled:

- o The MT-bit MUST be set in Hello packets and SHOULD be set in Database Description packets.

- o The router will only accept a Hello packet if the MT-bit is set (see Section 4.3).

When DefaultExclusionCapability is set to enabled, a router is said to be operating in DefaultExclusionCapability mode.

4.3. Adjacency Formation with Link Exclusion Capability

In order to have a smooth transition from a non-MT area to an MT-area, an MT router with DefaultExclusionCapability disabled will form adjacencies with non-MT routers and will include all links as part of the default topology.

A link may cease participating in the default topology if DefaultExclusionCapability is set to enabled. In this state, a router will only form adjacency with routers that set the MT-bit in their Hello packets. This will ensure that all routers have DefaultExclusionCapability enabled before the default topology can be disabled on a link.

Receiving OSPF Hello packets as defined in Section 10.5 of [OSPF] is modified as follows:

- o If the DefaultExclusionCapability in the Area Data structure is set to enabled, Hello packets are discarded if the received packet does not have the MT-bit set in the Header Options.

Receiving OSPF Database Description packets as defined in Section 10.6 of [OSPF] is unchanged. While packet options are validated in Hello packets, the only option checking performed for Database Description packets is ensuring that the options do not change during the database exchange process.

4.4. OSPF Control Packets Transmission over Excluded Links

If DefaultExclusionCapability is enabled, the default topology can be disabled on an interface. Disabling the default topology on an interface does not impact the installation of connected routes for the interface in the default topology. It only affects what a router advertises in its Router-LSA.

This allows OSPF control packets to be sent and received over an interface even if the default topology is disabled on the interface.

4.5. OSPF LSA Advertisement and SPF Computation for Excluded Links

When DefaultExclusionCapability is enabled and the link does not participate in the default topology, the MT-ID#0 metric is not advertised. The link’s TOS 0 metric is ignored during the default topology SPF computation.

When DefaultExclusionCapability is enabled and a link participates in the default topology, MT-ID#0 metric is used to advertise the metric associated with the default topology. The link’s TOS 0 metric is ignored during the default topology SPF computation.

Independent of the DefaultExclusionCapability, the TOS 0 metric is used for Summary-LSAs and Type-5 and Type-7 AS-external-LSAs.

- o If the prefix or router does not exist in the default topology, the TOS 0 metric is set to infinity (0xFFFFFF).
- o If the prefix or router exists in the default topology, the TOS 0 metric is used to advertise the metric in the default topology.

During the summary and external prefix calculation for the default topology, the TOS 0 metric is used for Summary-LSAs and Type-5 and Type-7 AS-external-LSAs.

5. Interoperability between MT-Capable and Non-MT-Capable Routers

The default metric field is mandatory in all LSAs (even when the metric value is 0). Even when a link or prefix does not exist in the default topology, a non-MT router will consider the zero value in the metric field as a valid metric and consider the link or prefix as part of the default topology.

In order to prevent the above problem, an MT-capable router will include all links as part of the default topology. If links need to be removed from the default topology, an MT-capable router must be configured in DefaultExclusionCapability mode. In this mode, routers

will ensure that all other routers in the area are in the DefaultExclusionCapability mode before considering the MT-ID#0 metric in the SPF calculation. Only then can the TOS 0 metric field in Router-LSAs be safely ignored during the default topology SPF computation.

Note that for any prefix or router to become reachable in a certain topology, a contiguous path inside that topology must exist between the calculating router and the destination prefix or router.

5.1. Demand Circuit Compatibility Considerations

A change to an area's DefaultExclusionCapability requires additional processing for area neighbors that are suppressing Hello packets as specified in "Extending OSPF to Support Demand Circuits" [DEMAND]. When the DefaultExclusionCapability for an area is changed, Hello suppression must be disabled for these neighbors for a period of RouterDeadInterval seconds. This implies that Hello packets are sent with the DC-bit clear as specified in Section 3.2.1 of [DEMAND] during this period. After RouterDeadInterval seconds, either the adjacency will be taken down due to rejection of Hello packets with a conflicting MT-bit or Hello suppression will be renegotiated.

6. Migration from Non-MT-Area to MT-Area

Introducing MT-OSPF into a network can be done gradually to allow MT routers and non-MT routers to participate in the default topology while MT routers participate in other topologies.

If there is a requirement to exclude some links from the default topology in an area, all routers in the area MUST be in DefaultExclusionCapability mode. In this section, we describe the migration steps to consider while transitioning from a non-MT network to an MT network.

Consider a network with a backbone area and a set of non-backbone areas functioning in standard OSPF mode. We would like to migrate to an MT network either partially or completely.

1. As required, part of an area is upgraded to be MT capable. The MT routers will interact with non-MT routers in the default topology and participate in other topologies as required.

2. If a new non-backbone area is created for MT routers, it may be configured in DefaultExclusionCapability mode since there is no interaction required with non-MT routers. In this mode, the default topology can be excluded on links as required.

3. If there are several non-backbone areas where MT is being used, it is desirable that the backbone area first be upgraded to be MT capable so that inter-area routing is ensured for MT destinations in different areas.

4. Gradually, the whole network can be made MT capable.

Note that inter-area routing for the MT-area still depends on the backbone area. Therefore, if different areas configured for a given topology need to communicate, the backbone area also needs to be configured for this topology.

7. MT Network Management Considerations

When multiple OSPF topologies exist within a domain, some of the routers can be configured to participate in a subset of the MTs in the network. This section discusses some of the options we have to enable operations or the network management stations to access those routers.

7.1. Create Dedicated Management Topology to Include All the Nodes

This approach is to set up a dedicated management topology or 'in-band' management topology. This 'mgmt' topology will include all the routers need to be managed. The computed routes in the topology will be installed into the 'mgmt' Routing Information Base (RIB). In the condition of the 'mgmt' topology uses a set of non-overlapping address space with the default topology, those 'mgmt' routes can also be optionally installed into the default RIB. The advantages of duplicate 'mgmt' routes in both RIBs include: the network management utilities on the system do not have to be modified to use specific RIB other than the default RIB; the 'mgmt' topology can share the same link with the default topology if so designed.

7.2. Extend the Default Topology to All the Nodes

Even in the case in which default topology is not used on some of the nodes in the IP forwarding, we may want to extend the default topology to those nodes for the purpose of network management. Operators SHOULD set a high cost on the links that belong to the extended portion of the default topology. This way, the IP data traffic will not be forwarded through those nodes during network topology changes.

8. Security Considerations

This document does not raise any security issues that are not already covered in [OSPF].

9. IANA Considerations

The T-bit as defined in [TOS-OSPF] for a router's TOS capability is redefined as the MT-bit in this document. IANA has assigned the MT-bit as defined in Section 4.1.

Similarly, the TOS field for Router-LSAs, Summary-LSAs, and Type-5 and Type-7 AS-external-LSAs, as defined in [OSPF], is redefined as MT-ID in Section 3.7.

IANA created a new registry, "OSPF Multi-Topology ID Values", with the assignments and registration policies listed in Section 3.7 of this document.

10. References

10.1. Normative References

[DEMAND] Moy, J., "Extending OSPF to Support Demand Circuits", RFC 1793, April 1995.

[NSSA] Murphy, P., "The OSPF Not-So-Stubby Area (NSSA) Option", RFC 3101, January 2003.

[OSPF] Moy, J., "OSPF Version 2", RFC 2328, April 1998.

[RFC-KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", RFC 3692, January 2004.

[TOS-OSPF] Moy, J., "OSPF Version 2", RFC 1583, March 1994.

10.2. Informative References

[M-ISIS] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in IS-IS", Work in Progress, October 2005.

[STUB] Retana, A., Nguyen, L., White, R., Zinin, A., and D. McPherson, "OSPF Stub Router Advertisement", RFC 3137, June 2001.

Appendix A. Acknowledgments

The authors would like to thank Scott Sturgess, Alvaro Retana, David Kushi, Yakov Rekhter, Tony Przygienda, and Naiming Shen for their comments on the document. Special thanks to Acee Lindem for editing and to Tom Henderson for an extensive review during the OSPF Working Group last call.

Appendix B. OSPF Data Formats

LSA content defined in [OSPF] is modified to introduce the MT-ID.

B.1. Router-LSAs

Router-LSAs are the Type 1 LSAs. Each router in an area originates a router-LSA. The LSA describes the state and cost of the router’s links (i.e., interfaces) to the area. All of the router’s links to the area must be described in a single router-LSA. For details concerning the construction of router-LSAs, see Section 12.4.1 of [OSPF].

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |            LS age             |     Options   |       1       |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Link State ID                          |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                     Advertising Router                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                     LS sequence number                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |         LS checksum           |             length            |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |*|*|*|N|W|V|E|B|        0      |            # links            |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                          Link ID                              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         Link Data                             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |     Type      |     # MT-ID   |            metric             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |     MT-ID     |       0       |          MT-ID  metric        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                              ...                              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |     MT-ID     |       0       |          MT-ID  metric        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                          Link ID                              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         Link Data                             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                              ...                              |
```

Figure 2: Router-LSA Format

B.2. Network-LSAs

Network-LSAs are the Type 2 LSAs. A network-LSA is originated for each broadcast and Non-Broadcast Multi-Access (NBMA) network in the area that supports two or more routers. The network-LSA is originated by the network's Designated Router. The LSA describes all routers attached to the network, including the Designated Router itself. The LSA's Link State ID field lists the IP interface address of the Designated Router.

The distance from the network to all attached routers is zero. This is why metric fields need not be specified in the network-LSA. For details concerning the construction of network-LSAs, see Section 12.4.2 of [OSPF].

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            LS age             |      Options  |      2        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Link State ID                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Advertising Router                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     LS sequence number                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         LS checksum           |             length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Network Mask                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Attached Router                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
```

Figure 3: Network-LSA Format

Note that network-LSA does not contain any MT-ID fields as the cost of the network to the attached routers is 0 and DR is shared by all topologies.

B.3. Summary-LSAs

Summary-LSAs are the Type 3 and 4 LSAs. These LSAs are originated by area border routers. Summary-LSAs describe inter-area destinations. For details concerning the construction of summary-LSAs, see Section 12.4.3 of [OSPF].

Type 3 summary-LSAs are used when the destination is an IP network. In this case the LSA's Link State ID field is an IP network number (if necessary, the Link State ID can also have one or more of the network's "host" bits set; see Appendix E of [OSPF] for details). When the destination is an AS boundary router, a Type 4 summary-LSA is used, and the Link State ID field is the AS boundary router's OSPF Router ID. (To see why it is necessary to advertise the location of each ASBR, consult Section 16.4 of [OSPF].) Other than the difference in the Link State ID field, the format of Type 3 and 4 summary-LSAs is identical.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            LS age             |     Options   |    3 or 4     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Link State ID                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Advertising Router                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     LS sequence number                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         LS checksum           |             length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Network Mask                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       0       |                  metric                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     MT-ID     |                MT-ID  metric                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     MT-ID     |                MT-ID  metric                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: Summary-LSA Format

B.4. AS-external-LSAs

AS-external-LSAs are the Type 5 LSAs. These LSAs are originated by AS boundary routers, and describe destinations external to the AS. For details concerning the construction of AS-external-LSAs, see Section 12.4.3 of [OSPF].

AS-external-LSAs usually describe a particular external destination. For these LSAs, the Link State ID field specifies an IP network number (if necessary, the Link State ID can also have one or more of the network's "host" bits set; see Appendix E of [OSPF] for details). AS-external-LSAs are also used to describe a default route. Default routes are used when no specific route exists to the destination. When describing a default route, the Link State ID is always set to DefaultDestination (0.0.0.0) and the Network Mask is set to 0.0.0.0.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            LS age             |     Options   |      5        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Link State ID                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Advertising Router                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     LS sequence number                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         LS checksum           |             length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Network Mask                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|E|     0       |                  metric                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Forwarding address                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      External Route Tag                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|E|    MT-ID    |              MT-ID  metric                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Forwarding address                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      External Route Tag                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|E|    MT-ID    |              MT-ID  metric                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Forwarding address                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      External Route Tag                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: AS-External-LSA Format

B.5. Type-7 AS-external-LSAs

Type-7 AS-external-LSAs are originated by AS boundary routers local to an NSSA (Not-So-Stubby Area), and describe destinations external to the AS. The changes to Type-7 AS-external-LSAs are identical to those for AS-external-LSAs (Appendix A.4.5 of [OSPF]). For details concerning the construction of Type-7 AS-external-LSAs, see Section 2.4 of [NSSA].

Authors' Addresses

Peter Psenak
Cisco Systems
Mlynske Nivy 43
821 09
Bratislava
Slovakia

EMail: ppsenak@cisco.com

Sina Mirtorabi
Force10 Networks
1440 McCarthy Blvd
Milpitas, CA 95035
USA

EMail: sina@force10networks.com

Abhay Roy
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

EMail: akr@cisco.com

Liem Nguyen
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

EMail: lhnguyen@cisco.com

Padma Pillay-Esnault
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

EMail: ppe@cisco.com

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.